arXiv:1609.08066v2 [math.OC] 5 Oct 2016

Integer programs with a fixed radius L1-constraint are polynomial-time solvable

Anil Aswani and Yonatan Mintz

Abstract Integer programming problems with L1-constraints are ubiquitous. For instance, L1-constraints are found in multidimensional integer packing problems with a capacity constraint. This paper shows that nonlinear integer programs with an L1-constraint can be solved in a number of oracle steps that is polynomial in the dimension of the decision variable, for each fixed radius of the L1-constraint. When specialized to integer (possibly non-convex) quadratically-constrained quadratic programs, our result shows such problems have a time complexity that is polynomial in simultaneously both the dimension of the decision variables and number of constraints, for each fixed radius of the L1-constraint. We prove this result using a geometric argument that leverages ideas from stochastic process theory and from the theory of convex bodies in high-dimensional spaces. Next we develop a polynomial-time algorithm that takes specific instances of integer programming problems with convex constraints and generates a polynomial or exponential in dimension upper bound on the number of oracle operations required to solve the problem. As a side result, we provide an additive polynomial time approximation scheme (PTAS) for continuous optimization of Lipschitz functions subject to Lipschitz constraints intersected with an L1-constraint.

Keywords nonlinear integer programming · approximation · covering number

1 Introduction.

Though it is well known that integer programming is NP-complete [1, 11, 14, 20, 21], a number of results use the notion of fixed-parameter complexity [8, 19] to provide a finer resolution (versus the problem classes of P or NP-complete) for characterizing

This work was supported in part by NSF Award CMMI-1450963 and the Philippine-California Advanced Research Institutes (PCARI).

A. Aswani and Y. Mintz
Industrial Engineering & Operations Research, University of California, Berkeley, CA, USA
E-mail: aaswani@berkeley.edu,E-mail: ymintz@berkeley.edu

complexity. Lenstra (and refinements by subsequent authors) [3,9,10,13,18] showed that the integer linear program (ILP) given by $\min\{c'x \mid Ax \leq b, x \in \mathbb{Z}^n\}$ with $c \in \mathbb{Q}^n$, $b \in \mathbb{Z}^m$, and $A \in \mathbb{Z}^{m\times n}$ can be: (1) solved in polynomial time in the number of constraints m when the dimension n is fixed, and (2) solved in polynomial time in dimension n when the number of constraints m is fixed. A related result [20] shows the ILP given by $\min\{c'x \mid Ax = b, x \in \mathbb{Z}^n_+\}$ with $c \in \mathbb{Q}^n$, $b \in \mathbb{Z}^m$, and $A \in \mathbb{Z}^{m\times n}$ can be solved in psuedopolynomial time when the number of constraints m is fixed. More recently, it has been shown that integer polynomial optimization $\min\{p(x) \mid Ax \leq b, x \in \mathbb{Z}^n\}$, where $p(x)$ is a polynomial satisfying certain properties, can be approximately [6,12] or exactly [7] solved in polynomial time for fixed dimension n.

However, little is known about which other classes of integer programs can be solved in fixed-parameter polynomial time. This paper proves that nonlinear integer programs of the form

$$\min\{f(x) \mid g(x) \leq 0, \|x\|_1 \leq \lambda, x \in \mathbb{Z}^n\} \tag{1}$$

can be solved in polynomial-time oracle complexity for fixed values of $\lambda \in \mathbb{R}_+$, where $\|\cdot\|_1$ is the usual ℓ_1-norm. Such constraints are common in integer programs: For instance, capacity constraints like $\{x : x \geq 0, \sum_i x_i \leq \lambda\}$ are equivalent to $\{x : x \geq 0, \|x\|_1 \leq \lambda\}$. If the functions f, g can be computed in polynomial time, such as linear or (possibly non-convex) quadratic functions, then our result implies (1) can be solved in polynomial (in simultaneously both the dimension n and number of constraints) time for fixed λ. This is in contrast to past algorithms [3,6,7,9,10, 12,13,18] that only apply to integer programs with polynomial objective and linear constraints and (depending on the algorithm) are polynomial in either dimension (for fixed number of constraints) or number of constraints (for fixed dimension).

Section 2 covers preliminaries, and Section 3 develops our algorithm to solve (1). We prove the number of integer-valued vectors in the scaled ℓ_1-ball is upper-bounded by a polynomial in n using stochastic process theory, and this new result is used to show that our algorithm has polynomial complexity when applied to problems with (weighted) L1-constraints. Next, Section 4 gives an algorithm to compute a polynomial or exponential in n upper bound on the number of oracle operations required to solve a given integer program $\min\{f(x) \mid g(x) \leq 0, x \in \mathbb{Z}^n\}$ for when: g is a convex function that can be optimized in polynomial time, and the continuous relaxation of the feasible set $\{x : g(x) \leq 0\}$ is bounded. Section 5 presents a side result in which we provide an additive polynomial time approximation schemes (PTAS) for optimizing

$$\min\left\{ f(x) \;\middle|\; g(x) \leq 0, \|x\|_1 \leq \lambda, x \in \mathbb{R}^n \right\}, \tag{2}$$

where f, g are Lipschitz: $|f(x) - f(y)| \leq \kappa\|x-y\|_\infty$ and $|g_i(x) - g_i(y)| \leq \kappa\|x-y\|_\infty$ for all i and some $\kappa \in \mathbb{R}_+$. This generalizes results that give a PTAS for optimization of Lipschitz functions [4] or polynomials [5,15] over the continuous unit simplex.

2 Preliminaries

Let $\mathbb{R}^n$, $\mathbb{Q}^n$, $\mathbb{Z}^n$, and $\mathbb{N}^n$ be the set of n-dimensional real-, rational-, integer-, and natural number-valued vectors, respectively, where the natural numbers are $\mathbb{N} = \{0, 1, \ldots\}$.

The notation $[r]$ defines the set $[r] := \{1,2,\ldots,r\}$, and $\lfloor r \rfloor$ is the largest integer smaller than r. The round(o) function rounds each component of the vector o to a nearest integer, and $\mathbf{1}(A)$ is an indicator function that is 1 if A is true, and 0 if A is false.

Recall the usual inner product $\langle g,x\rangle = \sum_j g_j x_j$, and let $\circ$ be the Hadamard or elementwise product operator, such that $z = x \circ y \iff z_i = x_i \cdot y_i,\ \forall i \in [n]$. We use the ℓ_p norm notation: $\|x\|_1 = \sum_j |x_j|$ for the ℓ_1-norm, $\|x\|_\infty = \max_j |x_j|$ for the ℓ_∞-norm, and $\|x\|_2 = \langle x,x\rangle^{1/2}$ for the ℓ_2-norm. Define $\mathcal{B}_1 = \{x : \|x\|_1 \le 1\}$, $\mathcal{B}_2 = \{x : \|x\|_2 \le 1\}$, and $\mathcal{B}_\infty = \{x : \|x\|_\infty \le 1\}$ to be the unit ℓ_1-, ℓ_2-, and ℓ_∞-balls centered at the origin, respectively. For any $\lambda \in \mathbb{R}_+$ and set $K \subset \mathbb{R}^n$, the scaled set is $\lambda K := \{\lambda x : x \in K\}$. For instance, $\lambda\mathcal{B}_1 = \{x : \|x\|_1 \le \lambda\}$, $\lambda\mathcal{B}_2 = \{x : \|x\|_2 \le \lambda\}$, and $\lambda\mathcal{B}_\infty = \{x : \|x\|_\infty \le \lambda\}$.

If $K, D \subset \mathbb{R}^n$ are convex sets, then the covering number of K by copies of D is defined as the quantity

$$N(K,D) = \min_V \left\{ \#V : \bigcup_{v\in V} \left(D \oplus v\right) \supseteq K \right\}, \tag{3}$$

where $\#V$ is the cardinality of the set V, and $\oplus$ denotes the Minkowski summation operator defined as $A \oplus B = \{a+b : a \in A, b \in B\}$. A related value is the packing number of K by copies of D, which is defined as the quantity

$$P(K,D) = \max_V \Big\{ \#V : \bigcup_{v\in V} \left(D \oplus v\right) \subseteq K \wedge \left(D \oplus v_i\right) \cap \left(D \oplus v_j\right) = \emptyset \text{ for all } v_i, v_j \in V \text{ with } v_i \neq v_j \Big\}. \tag{4}$$

A basic inequality [2] relating these two quantities is $P(K,2D) \le N(K,D) \le P(K,D)$.

3 Polynomial-Time Solvability with a Bounded L1-Constraint

This section shows (1) can be solved in polynomial time for fixed λ. We use stochastic process theory and the theory of convex bodies in high-dimensional spaces to characterize the complexity of scaled ℓ_p-balls, which we then use to upper bound the number of integers within scaled ℓ_p-balls centered at the origin. Surprisingly, we find that the number of integers in the scaled ℓ_1-ball is polynomial in dimension when the radius of the ball is fixed, which is in sharp contrast to the number of integers within the scaled ℓ_∞-ball. We use this result to design an algorithm to solve (1), and we prove its polynomial-time complexity. Our algorithm is then generalized to integer programs with a weighted L1-constraint.

3.1 Characterizing the Integer Complexity of Lp-Balls

Covering numbers provide one useful measure of the complexity of a set in Euclidean space. Exactly determining covering numbers is difficult, but fortunately stochastic process theory provides several approaches for bounding the covering number. One approach is Sudakov's minoration [17,22,24], which relates the covering number of a symmetric convex set K to its Gaussian mean width:

Proposition 1 (Sudakov's Minoration [22]) *Let $K \subset \mathbb{R}^n$ be a symmetric convex set, and recall that $N(K, r\mathcal{B}_2)$ is the covering number of K by ℓ_2-balls with radius r. We have the bound*

$$\sqrt{\log N(K, r\mathcal{B}_2)} \leq \frac{1}{2r}\mathbb{E}\Big(\max_{x\in K}\langle g, x\rangle\Big), \tag{5}$$

where $g \in \mathbb{R}^n$ is a vector whose entries are iid Gaussian random variables with zero mean and unit variance.

Another useful approach is a basic volume-based inequality [24] for bounding the covering number; however, these bounds can be loose.

Proposition 2 (Volume Estimate of Covering Number [24]) *For any symmetric convex sets $K, D \subset \mathbb{R}^n$ we have*

$$\frac{\text{vol}(K)}{\text{vol}(D)} \leq N(K, D) \leq \frac{\text{vol}(K \oplus \frac{1}{2}D)}{\text{vol}(D)}. \tag{6}$$

When $K/\lambda = D/r$, this simplifies to $(\frac{\lambda}{r})^n \leq N(\lambda K, rK) \leq (2 + \frac{\lambda}{r})^n$.

The above approaches for bounding the covering number of a symmetric convex set can be used to derive bounds for the ℓ_1-, ℓ_2-, and ℓ_∞-balls. Asymptotic bounds are found in [24], but we need non-asymptotic bounds for our purposes.

Proposition 3 *Recall that $\lambda\mathcal{B}_1$, $\lambda\mathcal{B}_2$, and $\lambda\mathcal{B}_\infty$ are the ℓ_1-, ℓ_2-, and ℓ_∞-balls centered at the origin and with radius λ. If $r \leq \lambda$, then we have*

$$\begin{aligned} 2n \leq \quad & N(\lambda\mathcal{B}_1, r\mathcal{B}_\infty) \leq \quad n^{(\lambda/\sqrt{2}r)^2} \\ 2n \leq \quad & N(\lambda\mathcal{B}_2, r\mathcal{B}_\infty) \leq \quad (2 + \tfrac{\lambda}{r})^n \\ (\tfrac{\lambda}{r})^n \leq \quad & N(\lambda\mathcal{B}_\infty, r\mathcal{B}_\infty) \leq \quad (2 + \tfrac{\lambda}{r})^n \end{aligned} \tag{7}$$

The upper bounds hold unconditionally, that is for all $\lambda \geq 0$.

Proof Since $\lambda\mathcal{B}_1$ is symmetric and convex, we can upper bound its covering number using the Sudakov minoration. The first step is to compute the Gaussian mean width, which is achieved by noting that Hölder's inequality and the symmetry of $\lambda\mathcal{B}_1$ give

$$\mathbb{E}\Big(\max_{x\in\lambda\mathcal{B}_1}\langle g, x\rangle\Big) \leq \mathbb{E}\Big(\max_{x\in\lambda\mathcal{B}_1}\|x\|_1 \cdot \max_j |g_j|\Big) \leq \lambda\sqrt{2\log n}, \tag{8}$$

where we have used the elementary bound $\max_j |g_j| \leq \sqrt{2\log n}$. Combining this with Sudakov's minoration (5) gives

$$\sqrt{\log N(\lambda\mathcal{B}_1, r\mathcal{B}_2)} \leq \lambda\sqrt{\log n}/\sqrt{2}r, \tag{9}$$

which simplifies to $N(\lambda\mathcal{B}_1, r\mathcal{B}_2) \leq n^{(\lambda/\sqrt{2}r)^2}$. The upper bound follows by noting $N(\lambda\mathcal{B}_1, r\mathcal{B}_\infty) \leq N(\lambda\mathcal{B}_1, r\mathcal{B}_2)$ since $\mathcal{B}_2 \subseteq \mathcal{B}_\infty$. The lower bound is given in [24], and the argument is that one copy of $r\mathcal{B}_2$ is needed for each vertex of $\lambda\mathcal{B}_1$.

Bounds for the covering number of $\lambda\mathcal{B}_\infty$ follow from the simplified expression in Proposition 2 for the situation where $K/\lambda = D/r = \mathcal{B}_\infty$.

An upper bound for the covering number of $\lambda\mathcal{B}_2$ follows by noting $N(\lambda\mathcal{B}_2, r\mathcal{B}_\infty) \leq N(\lambda\mathcal{B}_\infty, r\mathcal{B}_\infty)$ since $\mathcal{B}_2 \subseteq \mathcal{B}_\infty$, and a lower bound follows by noting $N(\lambda\mathcal{B}_2, r\mathcal{B}_\infty) \geq N(\lambda\mathcal{B}_1, r\mathcal{B}_\infty)$ since $\mathcal{B}_2 \supseteq \mathcal{B}_1$. □

Remark 1 This says the number of ℓ_∞-balls we need to cover an ℓ_1-ball is polynomial in n for fixed λ/r. This is significant because in general (e.g., the ℓ_∞-balls) covering an n-dimensional convex body requires an exponential in n number of balls.

The reason for our interest in covering numbers is that they can be used to count the number of integers within a convex set. In particular, we have the bounds:

Theorem 1 *Let $K \subset \mathbb{R}^n$ be a convex set. For any $\delta \in (0,2)$ we have that*

$$N(K, \tfrac{1}{2-\delta}\mathcal{B}_\infty) \leq \#(K \cap \mathbb{Z}^n) \leq N(K, \tfrac{1}{2+\delta}\mathcal{B}_\infty). \tag{10}$$

Proof The upper bound follows if we can show that two integer-valued vectors cannot lie within the same $\frac{1}{2+\delta}\mathcal{B}_\infty$ covering-ball. To prove this, suppose the opposite is true. Then there exists some point o such that u, v are two integer-valued vectors with $u \neq v$ and $u, v \in \{x : \|x-o\|_\infty \leq \frac{1}{2+\delta}\}$. Using the triangle inequality gives $\|u-v\|_\infty \leq \|u-o\|_\infty + \|v-o\|_\infty \leq \frac{2}{2+\delta} < 1$ which is a contradiction because $\|u-v\|_\infty \geq 1$.

The lower bound follows by using the basic inequality $N(K,D) \leq P(K,D)$ and noting that a packing ball $\frac{1}{2-\delta}\mathcal{B}_\infty$ must contain at least one integer, since if o is the center of the packing ball then $\|o - \text{round}(o)\|_\infty \leq \frac{1}{2} < \frac{1}{2-\delta}$ by definition; meaning $\text{round}(o)$ is an integer-valued vector in the packing ball $\frac{1}{2-\delta}\mathcal{B}_\infty$ centered at o. □

Remark 2 A simplified set of bounds implied by the above theorem are

$$N(K, \mathcal{B}_\infty) \leq \#(K \cap \mathbb{Z}^n) \leq N(K, \tfrac{1}{4}\mathcal{B}_\infty). \tag{11}$$

The above result can be used to bound the number of integer-valued vectors within the ℓ_1-, ℓ_2-, and ℓ_∞-balls.

Corollary 1 *Recall that $\lambda\mathcal{B}_1$, $\lambda\mathcal{B}_2$, and $\lambda\mathcal{B}_\infty$ are the ℓ_1-, ℓ_2-, and ℓ_∞-balls centered at the origin and with radius λ. If $\lambda \geq 1$, then for any $\delta \in (0,1)$ we have that*

$$\begin{aligned} 2n \leq \quad & \#(\lambda\mathcal{B}_1 \cap \mathbb{Z}^n) \leq \quad n^{((2+\delta)\lfloor\lambda\rfloor)^2/2} \\ 2n \leq \quad & \#(\lambda\mathcal{B}_2 \cap \mathbb{Z}^n) \leq \quad (1+2\lfloor\lambda\rfloor)^n \\ (1+2\lfloor\lambda\rfloor)^n = \quad & \#(\lambda\mathcal{B}_\infty \cap \mathbb{Z}^n) \end{aligned} \tag{12}$$

The upper bounds hold unconditionally, that is for all $\lambda \geq 0$.

Proof If $x \in \mathbb{Z}^n$, then $\|x\|_1 \leq \lambda$ implies $\|x\|_1 \leq \lfloor\lambda\rfloor$, which means $(\lambda\mathcal{B}_1 \cap \mathbb{Z}^n) = (\lfloor\lambda\rfloor\mathcal{B}_1 \cap \mathbb{Z}^n)$. The bounds for $\#(\lambda\mathcal{B}_1 \cap \mathbb{Z}^n)$ then follow by Theorem 1 and Proposition 3. Similarly, if $x \in \mathbb{Z}^n$, then $\|x\|_\infty \leq \lambda$ implies $\|x\|_\infty \leq \lfloor\lambda\rfloor$, which means $(\lambda\mathcal{B}_\infty \cap \mathbb{Z}^n) = (\lfloor\lambda\rfloor\mathcal{B}_\infty \cap \mathbb{Z}^n)$. However, we have $\#(\lfloor\lambda\rfloor\mathcal{B}_\infty \cap \mathbb{Z}^n) = (1+2\lfloor\lambda\rfloor)^n$ since each edge of the cube $\lfloor\lambda\rfloor\mathcal{B}_\infty$ contains $(1+2\lfloor\lambda\rfloor)$ integer points. Bounds for $\#(\lambda\mathcal{B}_2 \cap \mathbb{Z}^n)$ follow by noting $\mathcal{B}_1 \subseteq \mathcal{B}_2 \subseteq \mathcal{B}_\infty$. □

Remark 3 The relevant outcome for our purposes is that $\#(\lambda\mathcal{B}_1 \cap \mathbb{Z}^n)$ is polynomial in n when λ is fixed. As is well known, the above result says that $\#(\lambda\mathcal{B}_\infty \cap \mathbb{Z}^n)$ is exponential in n. Our bounds for the ℓ_2-ball are ambiguous, though we conjecture that $\#(\lambda\mathcal{B}_2 \cap \mathbb{Z}^n)$ is exponential in n.

Remark 4 A simplified set of bounds for $\lambda \geq 1$ implied by the above theorem are

$$2n \leq \#(\lambda\mathcal{B}_1 \cap \mathbb{Z}^n) \leq n^{4\lambda^2}. \tag{13}$$

Algorithm 1 Solve L1-constrained integer programs

$f^* := +\infty$
for all $u \in M_n^{\lfloor\lambda\rfloor+1}$ **do**
  $v := \varphi(u)$
  for all $i \in [n]$ **do**
    $b_i := \mathbf{1}(v_i \neq 0)$
  end for
  for all $s \in \text{SignPerm}(b)$ **do**
    $x := v \circ s$
    if $f(x) < f^*$ and $g(x) \leq 0$ **then**
      $x^* = x$
      $f^* = f(x)$
    end if
  end for
end for

3.2 Algorithm for Solving L1-Constrained Integer Programs

Let M_k^N be the set of vectors with values corresponding to the k-multisets of set $[N]$ (bags with k values chosen with replacement from the set $[N]$), such that for all $u \in M_k^N$ we have $u_i \leq u_{i+1}$ $\forall i \in [k-1]$. Standard algorithms can generate each single combination in $O(1)$ oracle time, $O(\log N)$ arithmetic time, and $O(k)$ memory [23]. Let $\text{SignPerm}(b)$ be the set of all sign $(+/-)$ permutations of the nonzero entries of the binary vector b. Standard algorithms can generate each single permutation in $O(1)$ oracle time, $O(\log k)$ arithmetic time, and $O(\sum_i b_i)$ memory [16].

Next, we define the function $\varphi : M_k^N \to \mathbb{N}^k$ such that for any $u \in M_k^N$ the function is given by $(\varphi(u))_1 = u_1 - 1$ and $(\varphi(u))_i = u_i - u_{i-1}$ for $i \in \{2,\ldots,k\}$. This function maps each $u \in M_k^N$ to a corresponding point in $(N-1)\mathcal{B}_1 \cap \mathbb{N}^k$. The below result shows this function is a bijection and is computable in linear time:

Lemma 1 *The function φ provides a bijection between M_k^N and $(N-1)\mathcal{B}_1 \cap \mathbb{N}^k$, and it is computable in $O(k)$ oracle steps and $O(k \log N)$ arithmetic steps.*

Proof Consider a vector $u \in M_k^N$. By construction each component of $\varphi(u)$ is a non-negative integer, and we have $\sum_{i=1}^k (\varphi(u))_i = u_1 - 1 + \sum_{i=2}^k (u_i - u_{i-1}) = u_k - 1 \leq N-1$. Thus $\varphi(u) \in (N-1)\mathcal{B}_1 \cap \mathbb{N}^k$. To show φ is bijective, observe that it can be written as $\varphi(u) = Mu - e_1$ where M is a lower bidiagonal (square) matrix with 1 on the main diagonal and -1 on the diagonal below, and e_1 is a vector whose first entry is 1 and the remaining entries are 0. The matrix M can be seen to have full rank, and so φ must be bijective. The computational complexity follows because we perform one subtraction operation per entry of the result vector, which has dimension k. □

Theorem 2 *Algorithm 1 solves (1) in $O(n^{((2+\delta)\lfloor\lambda\rfloor)^2/2+1})$ oracle steps, for any $\delta \in (0,1)$ and where computing f, g comprises a single oracle step.*

Proof Recall that $\#(\lambda\mathcal{B}_1 \cap \mathbb{Z}^n) \leq n^{((2+\delta)\lfloor\lambda\rfloor)^2/2}$ by Corollary 1, the sets $M_n^{\lfloor\lambda\rfloor+1}$, $\text{SignPerm}(h)$ are of finite cardinality and such that each single combination can be computed in $O(1)$ oracle time [16,23], and computing v requires $O(n)$ oracle time

by Lemma 1. Hence it suffices to show that Algorithm 1 will enumerate over every point $x \in \lambda\mathcal{B}_1 \cap \mathbb{Z}^n$ exactly once. Let $\mathcal{A}$ be the set of points that are enumerated by Algorithm 1. Suppose $x \in \mathcal{A}$ then by the second inner loop this must mean that $|x| \in \mathcal{A}$. Since $|x| = \varphi(u)$ for some $u \in M_n^{\lfloor\lambda\rfloor+1}$ this means by Lemma 1 that $\sum_{i=1}^n |x|_i \leq \lambda$. Hence $x \in \lambda\mathcal{B}_1 \cap \mathbb{Z}^n$ and $\mathcal{A} \subseteq \lambda\mathcal{B}_1 \cap \mathbb{Z}^n$. Now suppose $x \in \lambda\mathcal{B}_1 \cap \mathbb{Z}^n$, then $|x| \in \lambda\mathcal{B}_1 \cap \mathbb{N}^n$. By Lemma 1, φ is bijective: Therefore, there exists a unique $u \in M_n^{\lfloor\lambda\rfloor+1}$ such that $u = \varphi_0^{-1}(|x|)$. Since Algorithm 1 iterates over all $u \in M_n^{\lfloor\lambda\rfloor+1}$, $|x|$ must be generated by some iteration of the second inner loop. Hence $\lambda\mathcal{B}_1 \cap \mathbb{Z}^n \subseteq \mathcal{A}$, and so $\mathcal{A} = \lambda\mathcal{B}_1 \cap \mathbb{Z}^n$ since we have shown both set inclusions. Note that since φ is a bijection and each permutation vector in SignPerm is distinct, this means each point cannot be enumerated more then once. □

Remark 5 A simplified result implied by the above theorem is that the oracle complexity is $O(n^{4\lambda^2+1})$.

Corollary 2 *If f, g are computable in polynomial time, then Algorithm 1 solves (1) in polynomial time for fixed λ. When s is the most bits needed to represent the objective or a constraint, and $\delta \in (0,1)$; then some specific instances are:*

1. *Integer Linear Program (ILP): If $f = c'x$ and $g = Ax - b$, where $A \in \mathbb{Q}^{m\times n}$, $b \in \mathbb{Q}^m$, and $c \in \mathbb{Q}^n$; then Algorithm 1 solves (1) in $O(s^2mn^{((2+\delta)\lfloor\lambda\rfloor)^2/2+1})$ time.*
2. *(Non-convex) Integer Quadratic Program (IQP): If $f = x'Q'x + c'x$ and $g_i = x'A_ix + b_i'x + c_i$ for $i \in [m]$, where $A_i, Q \in \mathbb{Q}^{n\times n}$, $b_i, c \in \mathbb{Q}^n$, and $c_i \in \mathbb{Q}$; then Algorithm 1 solves (1) in $O(s^2mn^{((2+\delta)\lfloor\lambda\rfloor)^2/2+2})$ time.*
3. *(Non-convex) Integer Quadratically-Constrained Quadratic Program (IQCQP): If $f = x'Qx + c'x$ and $g_i = x'A_ix + b_i'x + c_i$ for $i \in [m]$, where $A_i, Q \in \mathbb{Q}^{n\times n}$, $b_i, c \in \mathbb{Q}^n$, and $c_i \in \mathbb{Q}$; then Algorithm 1 solves (1) in $O(s^2mn^{((2+\delta)\lfloor\lambda\rfloor)^2/2+2})$ time.*

Proof Using the result of Theorem 2 it suffices to compute the time complexity of the oracle for f, g to determine the overall time complexity of each instance, which will result in the form $O((P(n,m,s) + ns) \cdot n^{((2+\delta)\lfloor\lambda\rfloor)^2/2})$, where P is a polynomial of n, m, s and $O(ns)$ is the arithmetic complexity for generating each iteration by Lemma 1 and [16, 23]. For ILP, since the dot product of two vectors in $\mathbb{Q}^n$ can be computed in $O(s^2n)$ time and g is comprised of m constraints this means that the complexity f, g is of order $O(s^2mn)$ and hence the resulting complexity bound. For IQP, since resolving quadratic forms of matrices in $\mathbb{Q}^{n\times n}$ requires time complexity of $O(s^2n^2)$ we obtain that the complexity of evaluating f, g is of order $O(s^2mn^2)$. Likewise, for IQCQP the time complexity of resolving the quadratic constraints dominates the complexity of evaluating the objective function, hence the time complexity of resolving this oracle is $O(s^2mn^2)$. □

Remark 6 A simplified result from the above corollary is that the arithmetic complexity is: $O(s^2mn^{4\lambda^2+1})$ for ILP, and $O(s^2mn^{4\lambda^2+2})$ for IQP and IQCQP.

Algorithm 2 Bound running time of integer programs with convex constraints

for all $i \in [n]$ **do**
  $l_i := -\min\left\{\min_x \{x_i \mid g(x) \leq 0\}, 0\right\}$
  $u_i := \max\left\{\max_x \{x_i \mid g(x) \leq 0\}, 0\right\}$
end for
$\rho := \left\lfloor \max_x \left\{ \sum_i (s_i + t_i) \mid g(s-t) \leq 0, 0 \leq s \leq u, 0 \leq t \leq l \right\} \right\rfloor$
bnd $:= n^{((2+\delta)\rho)^2/2+1}$ for any $\delta \in (0,1)$

3.3 Algorithm for Solving Weighted L1-Constrained Integer Programs

The algorithm and analysis in the previous section easily generalize to the case of an integer program with a weighted L1-constraint:

$$\min\{ f(x) \mid g(x) \leq 0, \textstyle\sum_i w_i |x_i| \leq \lambda, x \in \mathbb{Z}^n \}, \quad \text{with } w_i > 0. \tag{14}$$

If we define the effective dimension $\overline{n} = \sum_i \mathbf{1}(w_i \leq \lambda)$, the effective radius $\mu = \lambda / (\min_i w_i)$, and the effective decision variable $y \in \mathbb{Z}^{\overline{n}}$ with a bijection between the components of y and the non-zero components of x (i.e., x_i such that $w_i \leq \lambda$); then we can rewrite this problem as

$$\min\{ f(My) \mid g(My) \leq 0, \textstyle\sum_i w_i |(My)_i| \leq \lambda, \|y\|_1 \leq \mu, y \in \mathbb{Z}^{\overline{n}} \}, \tag{15}$$

where M gives the bijection (i.e., $x = My$). We can then solve this problem by applying Algorithm 1, and similar results to Theorem 2 and Corollary 2 can be shown but with $\overline{n}$ and μ taking the place of n and λ.

4 Bounding the Running Time of a Given Integer Program

Here, we provide a polyomial-time algorithm to generate upper bounds on the number of oracle operations required to solve an integer program

$$\min\{ f(x) \mid g(x) \leq 0, x \in \mathbb{Z}^n \} \tag{16}$$

with a bounded feasible region and convex constraints $g(x)$ that can be optimized in polynomial time. Let $\mathcal{C} = \{x : g(x) \leq 0\}$ be the continuous relaxation for the feasible region of (16). Our algorithm is deterministic and provides bounds that are polynomial or exponential in dimension n. The trivial bound on complexity for when the feasible region lies within $\lambda \mathcal{B}_\infty$ is $(1 + 2\lfloor \lambda \rfloor)^n$ (see Corollary 1), and our algorithm can potentially provide polynomial in n bounds for specific problem instances.

Algorithm 2 describes our procedure. The intuition is that the algorithm finds a radius ρ such that $\rho \mathcal{B}_1$ covers $\mathcal{C}$, and then bounds the oracle complexity using Theorem 2. When the feasible set $\mathcal{C}$ is nonnegative (i.e., $\mathcal{C} \subset \mathbb{R}^n_+$), the minimum radius is readily computed by solving $\max_x \left\{ \sum_i x_i \;\middle|\; g(x_i) \leq 0 \right\}$. When $\mathcal{C}$ is general (i.e., $\mathcal{C} \subset \mathbb{R}^n$), directly finding the minimum radius is difficult in this case. So the algorithm takes a different approach: It performs a change of variables $x = u - v$ where $u, v \geq 0$, computes bounds on u, v, and converts (16) into another instance of (16) with a nonnegative feasible region. The correctness and polynomial-time complexity of Algorithm 2 is given by the following result:

Theorem 3 *Suppose $\mathcal{C} \subset \mathbb{R}^n$ is bounded, and that $g(x)$ can be optimized in polynomial time. Then Algorithm 2 runs in polynomial time, and the value* bnd *computed by the algorithm is an upper bound on the oracle complexity for solving (16).*

Proof This algorithm runs in polynomial time because it consists of solving $2n+1$ convex optimization problems with at most $2n$ variables, each of which can be solved in polynomial time by assumption on g and the linearity of the objective functions. Next we prove that the value bnd provides an upper bound on the oracle complexity. If $x = s - t$ with $s,t \geq 0$; then $\|x\|_1 = \|s-t\|_1 \leq \|s\|_1 + \|t\|_1 = \sum_i (s_i + t_i)$, where the last equality holds since $s,t \geq 0$. And l,u are constructed so that for all $x \in \mathcal{C}$: if $s = \max\{x,0\}$ and $t = \max\{-x,0\}$, then we have $x = s - t$ with $0 \leq s \leq u$ and $0 \leq t \leq l$. Hence we have $\max_x \left\{ \|x\|_1 \;\middle|\; g(x) \leq 0 \right\} \leq \max_x \left\{ \sum_i (s_i + t_i) \;\middle|\; g(s-t) \leq 0, 0 \leq s \leq u, 0 \leq t \leq l \right\}$. Thus $\mathcal{C} \cap \mathbb{Z}^n \subseteq \rho \mathcal{B}_1 \cap \mathbb{Z}^n$ for ρ as defined in the algorithm. This means we can solve (16) by using Algorithm 1 to solve (1) with $\rho \equiv \lambda$, and so the oracle complexity for solving (16) is bounded by the rate given in Theorem 2. □

Remark 7 When $\mathcal{C}$ lies in a single orthant, the value ρ is tight in the sense that $\max_x \left\{ \|x\|_1 \;\middle|\; g(x) \leq 0 \right\} = \max_x \left\{ \sum_i (s_i + t_i) \;\middle|\; g(s-t) \leq 0, 0 \leq s \leq u, 0 \leq t \leq l \right\}$.

Remark 8 Algorithm 2 can be modified to return simplified bounds by changing the last statement of the algorithm to "bnd $:= n^{4\rho^2+1}$".

The next proposition shows that our algorithm returns non-trivial bounds. More specifically, it returns polynomial bounds for some instances of (16) and exponential bounds for other instances of (16).

Proposition 4 *Suppose $\lambda \in \mathbb{N}$. If $\mathcal{C} \subseteq \lambda \mathcal{B}_1 \cap \mathbb{R}^n_+$, then Algorithm 2 returns a value* $\text{bnd} = O(n^{4\lambda^2+1})$ *that is polynomial in n. If $\mathcal{C} \supseteq \lambda \mathcal{B}_\infty \cap \mathbb{R}^n_+$, then Algorithm 2 returns a value* $\text{bnd} = \Omega(n^{2(n\lambda)^2+1}) = \Omega((1+2\lambda)^n)$ *that is exponential in n.*

Proof In both cases, $\mathcal{C}$ is in the non-negative orthant and so $\max_x \left\{ \|x\|_1 \;\middle|\; g(x) \leq 0 \right\} = \max_x \left\{ \sum_i (s_i + t_i) \;\middle|\; g(s-t) \leq 0, 0 \leq s \leq u, 0 \leq t \leq l \right\}$. The first case has $\mathcal{C} \subseteq \lambda \mathcal{B}_1 \cap \mathbb{R}^n_+$, and so $\max_x \left\{ \|x\|_1 \;\middle|\; g(x) \leq 0 \right\} \leq \max_x \left\{ \|x\|_1 \;\middle|\; x \in \lambda \mathcal{B}_1 \cap \mathbb{R}^n_+ \right\} = \lambda$. Thus the algorithm will compute $\rho \leq \lambda$, which means it will return bnd $\leq n^{((2+\delta)\lambda)^2/2+1}$ for any $\delta \in (0,1)$. The second case has $\mathcal{C} \supseteq \lambda \mathcal{B}_\infty \cap \mathbb{R}^n_+$, and so $\max_x \left\{ \|x\|_1 \;\middle|\; g(x) \leq 0 \right\} \geq \max_x \left\{ \|x\|_1 \;\middle|\; x \in \lambda \mathcal{B}_\infty \cap \mathbb{R}^n_+ \right\} = n\lambda$. Thus the algorithm will compute $\rho \geq n\lambda$, which means it will return bnd $\geq n^{((2+\delta)n\lambda)^2/2+1}$ for any $\delta \in (0,1)$. □

Remark 9 We have used simplified bounds, without $\delta \in (0,1)$, when stating the above proposition.

Remark 10 Algorithm 2 returns a hyper-exponential bound bnd $= \Omega(n^{2(n\lambda)^2+1})$ when $\mathcal{C} \supseteq \lambda \mathcal{B}_\infty \cap \mathbb{R}^n_+$, and we include the weaker (exponential in n) bound $\Omega((1+2\lambda)^n)$ to emphasize that our algorithm provides a bound that is consistent with Corollary 1.

Algorithm 3 PTAS for continuous optimization of Lipschitz problems

$f^* := +\infty$
for all $u \in M_n^{\lfloor \lambda\kappa/\varepsilon \rfloor + 1}$ **do**
  $v := \varphi(u)$
  for all $i \in [n]$ **do**
    $b_i := \mathbf{1}(v_i \neq 0)$
  end for
  for all $s \in \text{SignPerm}(b)$ **do**
    $x := \varepsilon/\kappa \cdot (v \circ s)$
    if $f(x) < f^*$ and $g(x) \leq \varepsilon$ **then**
      $x^* = x$
      $f^* = f(x)$
    end if
  end for
end for

5 PTAS for Optimizing Lipschitz Problems over the Scaled L1-Ball

In this section, we modify Algorithm 1 in order to develop an additive PTAS for (2). Let x^* be any minimizer of (2). Then we define an additive PTAS for (2) to be an algorithm that for fixed $\varepsilon > 0$ requires a polynomial in n number of oracle operations to compute a solution $\hat{x}$ such that $f(\hat{x}) - f(x^*) \leq \varepsilon$, $g(\hat{x}) \leq \varepsilon$, $\|\hat{x}\|_1 \leq \lambda$, and $\hat{x} \in \mathbb{R}^n$. Recall $\kappa \in \mathbb{R}_+$ is the Lipschitz constant in the sense: $|f(x) - f(y)| \leq \kappa\|x-y\|_\infty$ and $|g_i(x) - g_i(y)| \leq \kappa\|x-y\|_\infty$ for all i. Our Algorithm 3 finds such an $\hat{x}$ by enumerating over a set of points that forms an $(\varepsilon/\kappa)\mathcal{B}_\infty$ cover of the feasible region in polynomial time, and our final result formally proves this:

Theorem 4 *If $\varepsilon > 0$, and an optimal x^* exists for* (2)*; then Algorithm 3 is an additive PTAS for* (2) *with oracle time complexity $O(n^{((2+\delta)\lfloor\lambda\kappa/\varepsilon\rfloor)^2/2+1})$ for any $\delta \in (0,1)$.*

Proof Algorithm 3 is a modified version of Algorithm 1, and so the same argument from the proof of Theorem 2 implies $v \circ s$ enumerates over all points in $(\lambda\kappa/\varepsilon)\mathcal{B}_1 \cap \mathbb{Z}^n$. This means the x in Algorithm 3 are such that $\|x\|_1 = (\varepsilon/\kappa) \cdot \|v \circ s\|_1 = (\varepsilon/\kappa) \cdot \|v\|_1 \leq (\varepsilon/\kappa) \cdot \lfloor \lambda\kappa/\varepsilon \rfloor \leq \lambda$. Suppose these x are the centers of copies of $(\varepsilon/\kappa)\mathcal{B}_\infty$ that form a covering of $\lambda\mathcal{B}_1$, then there exists a covering ball containing x^*. Let $\tilde{x}$ be the center of this ball, and observe that $f(\tilde{x}) - f(x^*) \leq \kappa\|\tilde{x} - x^*\|_\infty \leq \varepsilon$ and similarly $g_i(\tilde{x}) \leq g_i(x^*) + |g_i(\tilde{x}) - g_i(x^*)| \leq \kappa\|\tilde{x} - x^*\|_\infty \leq \varepsilon$ for all i. This would mean that a solution $\hat{x}$ with the desired properties is returned by Algorithm 3. Furthermore, the same argument for Theorem 2 yields that the oracle run time complexity for Algorithm 3 is $O(n^{((2+\delta)\lfloor\lambda\kappa/\varepsilon\rfloor)^2/2+1})$.

Hence the result follows if we can show that the set of all points x generated by the algorithm are the centers of $(\varepsilon/\kappa)\mathcal{B}_\infty$ copies that form a covering of $\lambda\mathcal{B}_1$. Consider any $y \in \lambda\mathcal{B}_1$, and note $\lfloor \kappa|y|/\varepsilon \rfloor \in (\lambda\kappa/\varepsilon)\mathcal{B}_1 \cap \mathbb{Z}^n_+$. Thus Algorithm 3 must choose v such that $v = \lfloor \kappa|y|/\varepsilon \rfloor$ in some iteration by the same argument from the proof of Theorem 2. If we let $s = \text{sign}(y)$, then for this v, s the corresponding value of x chosen by the algorithm is $x = (\varepsilon/\kappa) \cdot (v \circ s) = (\varepsilon/\kappa) \cdot \lfloor \kappa|y|/\varepsilon \rfloor \circ s$. Thus $\left\|y - x\right\|_\infty = \left\|y - (\varepsilon/\kappa) \cdot \lfloor \kappa|y|/\varepsilon \rfloor \circ s\right\|_\infty = \left\||y| \circ s - (\varepsilon/\kappa) \cdot \lfloor \kappa|y|/\varepsilon \rfloor \circ s\right\|_\infty \leq \varepsilon/\kappa$. Restated, this argument shows that for any point $y \in \lambda\mathcal{B}_1$, Algorithm 3 generates an x such that

$\|y-x\|_\infty \leq \varepsilon/\kappa$. Thus the set of all points x generated by the algorithm are the centers of $(\varepsilon/\kappa)\mathcal{B}_\infty$ copies that form a covering of $\lambda\mathcal{B}_1$. □

Remark 11 A simplified result implied by the above theorem for $\varepsilon > 0$ is that the oracle complexity is $O(n^{4(\lambda\kappa/\varepsilon)^2+1})$

Remark 12 Algorithm 3 can be suitably modified, as in Section 3.3, to solve minimization problems with weighted L1-constraints.

6 Conclusion

Using a geometric argument based on stochastic process theory and the theory of convex bodies in high-dimensional spaces, we showed the number of integers within a scaled ℓ_1-ball is polynomial in dimension when the radius of the ball is fixed. This result was used to develop an algorithm that solves L1-constrained integer programs and has oracle complexity that is polynomial in dimension when the radius of the L1-constraint is fixed. Our result implies polynomial (in both the dimension of the decision variables and constraints) arithmetic time complexity for instances of integer programming with fixed radius L1-constraints and polynomial-time computable objective function and constraints (e.g., problems where the objective and constraints are linear or quadratic). Next we used these results to develop a polynomial-time algorithm that generates polynomial or exponential in dimension bounds on the oracle complexity for solving integer programs with convex constraints, and an additive PTAS for continuous optimization of problems with Lipschitz objective and Lipschitz constraints intersected with an L1-constraint.

References

1. Borosh, I., Treybig, L.B.: Bounds on positive integral solutions of linear diophantine equations. Proceedings of the American Mathematical Society **55**(2), 299–304 (1976)
2. Boucheron, S., Lugosi, G., Massart, P.: Concentration Inequalities: A Nonasymptotic Theory of Independence. OUP Oxford (2013)
3. Clarkson, K.L.: Las Vegas algorithms for linear and integer programming when the dimension is small. Journal of the ACM (JACM) **42**(2), 488–499 (1995)
4. De Klerk, E., Den Hertog, D., Elabwabi, G.: On the complexity of optimization over the standard simplex. European journal of operational research **191**(3), 773–785 (2008)
5. De Klerk, E., Laurent, M., Parrilo, P.A.: A PTAS for the minimization of polynomials of fixed degree over the simplex. Theoretical Computer Science **361**(2), 210–225 (2006)
6. De Loera, J.A., Hemmecke, R., Köppe, M., Weismantel, R.: Integer polynomial optimization in fixed dimension. Mathematics of Operations Research **31**(1), 147–153 (2006)
7. Del Pia, A., Hildebrand, R., Weismantel, R., Zemmer, K.: Minimizing cubic and homogeneous polynomials over integers in the plane. Mathematics of Operations Research **41**(2), 511–530 (2016)
8. Downey, R.G., Fellows, M.R.: Parameterized complexity. Springer Science & Business Media (1999)
9. Eisenbrand, F.: Fast integer programming in fixed dimension. In: European Symposium on Algorithms, pp. 196–207. Springer (2003)
10. Frank, A., Tardos, É.: An application of simultaneous diophantine approximation in combinatorial optimization. Combinatorica **7**(1), 49–65 (1987)
11. von zur Gathen, J., Sieveking, M.: A bound on solutions of linear integer equalities and inequalities. Proceedings of the American Mathematical Society **72**(1), 155–158 (1978)

12. Hildebrand, R., Weismantel, R., Zemmer, K.: An FPTAS for minimizing indefinite quadratic forms over integers in polyhedra. In: Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms, pp. 1715–1723. SIAM (2016)
13. Kannan, R.: Minkowski's convex body theorem and integer programming. Mathematics of operations research **12**(3), 415–440 (1987)
14. Kannan, R., Monma, C.L.: On the computational complexity of integer programming problems. In: Optimization and Operations Research, pp. 161–172. Springer (1978)
15. de Klerk, E., Laurent, M., Sun, Z.: An alternative proof of a PTAS for fixed-degree polynomial optimization over the simplex. Mathematical Programming **151**(2), 433–457 (2015)
16. Knuth, D.: The Art of Computer Programming, Volume 4A: Combinatorial Algorithms. pt. 1. Pearson Education (2014)
17. Ledoux, M., Talagrand, M.: Probability in Banach Spaces: Isoperimetry and Processes. A Series of Modern Surveys in Mathematics Series. Springer (1991)
18. Lenstra Jr, H.W.: Integer programming with a fixed number of variables. Mathematics of operations research **8**(4), 538–548 (1983)
19. Niedermeier, R.: Invitation to Fixed-Parameter Algorithms. Oxford Lecture Series in Mathematics and Its Applications. OUP Oxford (2006)
20. Papadimitriou, C.H.: On the complexity of integer programming. Journal of the ACM (JACM) **28**(4), 765–768 (1981)
21. Pia, A.D., Dey, S.S., Molinaro, M.: Mixed-integer quadratic programming is in NP. Mathematical Programming pp. 1–16 (2016)
22. Sudakov, V.: Gaussian random processes and solid angle measures in Hilbert space. Doklady Akademii Nauk SSSR **197**(1), 43 (1971)
23. Takaoka, T.: O(1) time algorithms for combinatorial generation by tree traversal. The Computer Journal **42**(5), 400–408 (1999)
24. Vershynin, R.: Lectures in geometric functional analysis. Tech. rep., University of Michigan (2009)